arXiv:1404.5510v1 [cs.DS] 22 Apr 2014

Serving Online Demands with Movable Centers

Abdolhamid Ghodselahi and Fabian Kuhn
Department of Computer Science, University of Freiburg
79110 Freiburg, Germany
{hghods,kuhn}@cs.uni-freiburg.de

Abstract

We study an online problem in which a set of mobile resources or centers have to be moved in order to optimally serve a set of demands. There is a set of n nodes and a set of k resources that are placed at some of the nodes. Each node can potentially host several resources and the resources can be moved between the nodes. At the nodes, demands arrive in an online fashion, and the cost for serving the demands is a function of the number of demands and resources at the different nodes. An online algorithm has to move resources in order to keep the service cost low, however as moving resources is expensive, the objective of an online algorithm is to minimize the total number of movements.

Specifically, we give a deterministic online algorithm, which for parameters $\alpha \geq 1$ and $\beta \geq 0$ guarantees that at all times, the service cost is within a multiplicative factor α and an additive term β of the optimal service cost and where also the movement cost is within a small multiplicative factor and an additive term of the optimal overall cost. Roughly, for $\alpha = 1$ and $\beta = \Omega(k)$, we show that the movement cost by time t is upper bounded by the optimal service cost at time t plus an additive term of order $O(k \log k)$. For $\alpha > 1$ and arbitrary β, we show that the movement cost by time t is only linear in $k \log k$ and logarithmic in the optimal service cost at time t. We also show that both bounds are almost asymptotically tight.

1 Introduction

Some of the most fundamental combinatorial optimization problems are of the following general type. There is a set V of nodes or sites and there is a set of demands or requests such that each demand is placed at some node in V. Further, there is a (possibly limited) set of resources or centers that have to be activated or placed at some nodes in V in order to serve the demands. Classic examples, e.g., include facility location and k-median [3, 5, 9, 15, 17]

In the present paper, we consider the following online problem of the above type. There is a set V of n nodes and there are k resources, where each resource has to be placed at one of the nodes. Further, there are demands that arrive at the nodes in an online fashion. We assume that any node can potentially host an arbitrary number of resources. Formally, the cost for serving the demands at each node v is given by a general cost function that depends on v, on the number of demands at node v, as well as on the number of resources placed at v. Generally, the more demands there are at some nodes, the more it costs to serve these demands. Further, if we place more resources at a given node, the cost for serving the demands at this node becomes smaller. A somewhat similar type of cost function has also been studied in the context of facility location in [16]. In [16], it is assumed that the cost of a facility increases as a function of the demands it needs to serve.

As the demands arrive one by one, we study online algorithms that possibly have to react whenever a new demand arrives. We assume that the resources are movable and the objective of an algorithm is to move around some of the resources in order to keep the overall service cost low. As moving a resource is an

expensive operation, we need to keep the number of movements as small as possible. An online algorithm has to be competitive with an optimal offline algorithm that can wait until all demands have arrived and just perform all the necessary resource movements at the very end.

We believe that the abstract problem studied in this paper might have applications in cases where resources are relatively large entities such that while they can be moved, doing this is a relatively costly operation, irrespective of, e.g., between which nodes a movement occurs. Think for example of a company that has several project teams. Moving a whole team to a new location is expensive, however depending on where new customers arrive, it might still be desirable to do. The cost for serving the customers at as certain location clearly depends (in a possibly non-linear way) on the number of project teams and on the number of customers at the location. Alternatively think of a distributed service that is offered on a large network such as the Internet. To offer the service, a provider might have a budget to place k servers at some nodes in the network. The best placement of servers depends on the distribution of the users of the distributed service and as the set of users might grow (or even change arbitrarily) over time, from time to time, we might have to move some of the servers, even though migrating a whole server might be a relatively costly thing to do.

In Section 4.1 we give a deterministic online algorithm that guarantees to bound the movement cost while keeping the overall service cost close to optimal at all times. More specifically, we fix two parameters $\alpha \geq 1$ and $\beta \geq 0$ and we require that given α and β, an algorithm has to guarantee that at all times, the overall service cost of the online algorithm is within a multiplicative factor α and an additive term β of the current optimal service cost. We show that for $\alpha = 1$, if $\beta = \Omega(k)$, the number of movements of our online algorithm by time t is at most $S_t^* + O(k \log k)$, where S_t^* denotes the optimal service cost at time t. For $\alpha > 1$, we show that for any $\beta \geq 0$, the movement cost of our algorithm is at most $\tilde{O}(k \log S_t^*)$, where the $\tilde{O}$ notation hides logarithmic factors in k and some model parameters. In addition, in Section 5, we prove a lower bound, which shows that both these bounds are best possible up to at most a logarithmic factor. For a formal and more general statement of the contributions of the present paper, we refer to Section 3.

Related Work: The most closely related of the classic problems are the facility location and the k-median problems.[1] Both problems have been widely studied in both computer science and operations research [3, 5, 6, 9, 15, 17]. When studying facility location, usually the position of the facilities is fixed. A variant where facilities are movable has been introduced as the mobile facility location in [7, 14] as a generalization of the standard k-median and facility location problems. The problem falls into a larger class of movement problems, which were introduced by Demaine et al. in [7]. In [1, 14], mobile facility location is modeled in such a way that the algorithm moves each facility and client to a point where in the final configuration, each client is at a node with some facility. The goal is to minimize the total movement cost of facilities and clients. We hope that our paper is also useful as a first step towards analyzing the movement cost of the online version of the mobile facility location problem. Ahmadian et al. [1] recently presented a $(3+\epsilon)$-factor approximation algorithm, matching the best-known approximation factor for the k-median problem [3]. An approximation-preserving reduction in [14] also shows that the mobile facility location indeed generalizes the k-median problem.

There exist various natural models in which the locations of demands are not known in advance, and a solution must be built or maintained gradually over time without any knowledge about future demands. The first algorithm for online facility location was introduced by Meyerson [20]. The paper presents a randomized algorithm with a competitive ratio of $O(\log n)$. Fotakis observed that with a simple modification of the same argument, the competitive ratio can be improved to $O\big(\frac{\log n}{\log \log n}\big)$ [12]. The paper also gives a

[1] Note that there are some significant differences between the problem studied in this paper on the one hand and facility location and k-median on the other hand. In facility location and k-median, having some distance function between nodes and facilities/centers is central. In our problem, service and also movement cost do not depend on any distances. However, unlike classically in facility location or k-center, we model the service cost by a general (possibly non-linear) cost function.

deterministic algorithm with the same competitive ratio and this is shown to be optimal. Further, an online algorithm that is computationally more efficient has been presented in [2].

Another online type of the problem, known as incremental facility location (and also k-median), is studied in [10]. In this variant of the problem, it is possible to close open facilities and assign all their demands to another open facility. Incremental algorithms were introduced in the context of clustering applications where merging clusters is quite natural [6]. A relaxed variant of incremental facility location is presented in [8] (also see [11]). In this variant, similarly as in our problem, the facilities can be moved to reduce the overall cost. However unlike in our paper, in [8], moving a facility is for free.

For a broader discussion of models and results on online and incremental facility location problems, we refer to the recent detailed survey in [13].

Finally, a classic online algorithm where one has to serve a set of online requests while trying to minimize the movement cost of the resources is the k-server problem [4, 18, 19]. While in our model all past requests have to be permanently served, in the k-server problem, each request only has to be served once when it arrives by moving some resource to the location of the request. Hence, the goal of the k-server problem is to minimize the total movement cost[2] needed to serve all requests. In our case, we at all times need to keep both the current service cost for serving all present request, as well as the overall movement cost as low as possible.

2 Problem Statement

We are given a set V of n nodes and there is a set of k resources. Further, there are demands $1, 2, \ldots$ that arrive one at a time. We assume that at time $t \geq 1$, demand t arrives at node $v(t) \in V$. For a node $v \in V$, let $d_{v,t}$ be the number of demands at node v after t demands have arrived, i.e., $d_{v,t} := |\{i \leq t : v(i) = v\}|$. In order to keep the overall service cost small, an algorithm can move the resources between the nodes (if necessary, for answering one new demand, we allow an algorithm to also move more than one resource). However throughout the execution, each of the k resources is always placed at one of the nodes $v \in V$. We define a *feasible solution* for an algorithm by integers $f_v \in \mathbb{N}_0$ for each $v \in V$ such that $\sum_{v \in V} f_v = k$. We describe such a feasible solution by a set F of pairs as follows:

$$F := \{(v, f_v) : v \in V\}. \tag{1}$$

We use F_0 to denote the initial configuration of the k resources. Further, for a given (deterministic) algorithm $\mathcal{A}$, for any $m \geq 1$, we let $F_m^{\mathcal{A}} = \left\{(v, f_{v,m}^{\mathcal{A}}) : v \in V\right\}$ be the solution of $\mathcal{A}$ after m resource movements (i.e., after m resource movements of $\mathcal{A}$, node v has $f_{v,m}^{\mathcal{A}}$ resources). Throughout the rest of the paper, we use $\tau_m^{\mathcal{A}}$ to denote the time of the m^{th} movement. In addition, we define $v_m^{src,\mathcal{A}}$ and $v_m^{dst,\mathcal{A}}$ to be the nodes involved in the m^{th} resource move, where we assume that $\mathcal{A}$ moves a resource from node $v_m^{src,\mathcal{A}}$ to $v_m^{dst,\mathcal{A}}$.

Service cost: We implicitly assume that if a node v has some resources, all demands at v are served by these resources (this also implies that the "assignment" of demands to resources can change over time). Depending on the number of resources and the number of demands at a node $v \in V$, the algorithm has to pay some *service cost* to serve the demands located at v. This service cost of node v is defined by a *service cost function* σ_v such that $\sigma_v(x, y) \geq 0$ is the cost for serving y demands if there are x resources at node v. For convenience, for $t \geq 1$, we also define $\sigma_{v,t}(x) := \sigma_v(x, d_{v,t})$ to be the service cost with x resources at node v at time t. For a feasible solution $F = \{(v, f_v) : v \in V\}$, we denote the total service cost at time t by $S_t(F)$:

$$S_t(F) := \sum_{v \in V} \sigma_{v,t}(f_v) = \sum_{v \in V} \sigma_v(f_v, d_{v,t}). \tag{2}$$

[2] In the special case of a uniform metric, one has to minimize the total number of movements and the problem is equivalent to the online paging problem.

For a given (deterministic) online algorithm $\mathcal{A}$, we denote the (feasible) solution at time t by $F^{\mathcal{A}}(t)$, where $F^{\mathcal{A}}(t) := \big\{(v, f_v^{\mathcal{A}}(t)) : v \in V\big\}$. Hence, $f_v^{\mathcal{A}}(t)$ denotes the number of resources at node v at time t (after $\mathcal{A}$ has reacted to the first t demands and immediately before demand $t+1$ arrives). The total service cost $S_t^{\mathcal{A}}$ of an algorithm $\mathcal{A}$ at time t (i.e., after t demands) can then be written as

$$S_t^{\mathcal{A}} := S_t\big(F^{\mathcal{A}}(t)\big) = \sum_{v \in V} \sigma_{v,t}\big(f_v^{\mathcal{A}}(t)\big) = \sum_{v \in V} \sigma_v\big(f_v^{\mathcal{A}}(t), d_{v,t}\big). \tag{3}$$

Movement cost: In addition, we define the *movement cost* $M_t^{\mathcal{A}}$ of an algorithm $\mathcal{A}$ to be the total number of resource movements by time t. Generally, for two feasible solutions, $F = \{(v, f_v) : v \in V\}$ and $F' = \{(v, f'_v) : v \in V\}$, we define a distance $\chi(F, F')$ between the two solutions as follows:

$$\chi(F, F') := \sum_{v \in V} \max\{0, f_v - f'_v\} = \frac{1}{2} \cdot \sum_{v \in V} |f_v - f'_v|. \tag{4}$$

The distance $\chi(F, F')$ is equal to the number of movements that are needed to get from configuration F to F' (or vice versa). Based on the definition of χ, we can also express the movement cost of an algorithm $\mathcal{A}$ as

$$M_t^{\mathcal{A}} = \sum_{i=1}^{t} \chi\big(F^{\mathcal{A}}(i-1), F^{\mathcal{A}}(i)\big). \tag{5}$$

Total cost: Our objective is to minimize the overall cost of the online algorithm, that is, we aim to keep $S_t^{\mathcal{A}} + M_t^{\mathcal{A}}$ as small as possible. To achieve this, our online algorithm keeps the service cost $S_t^{\mathcal{A}}$ within predefined bounds it and simultaneously tries to minimize the movement cost $M_t^{\mathcal{A}}$.

Optimal offline cost: An offline algorithm knows the set of demands in advance and can therefore move the resources after all demands have arrived. As a consequence, while the optimal service cost is potentially unbounded, the movement cost of an offline solution is at most k. In Appendix A, we show that the movement cost of an offline algorithm can also be accounted for by appropriately adjusting the service cost function σ. Formally, we therefore only charge an offline algorithm for its service cost. Let $\mathcal{F}^*(t)$ be the set of optimal solutions at time t (i.e., if the set of demands consists of exactly the first t demands). We use S_t^* to denote the optimal offline service cost at time t. We therefore have

$$S_t^* := \min_{F : F \text{ feasible}} S_t(F) \quad \text{and} \quad \mathcal{F}^*(t) := \{F : F \text{ is feasible} \wedge S_t(F) = S_t^*\}. \tag{6}$$

Service cost function properties: In order for the competitive analysis to work, the service cost function σ has to satisfy a number of natural properties. First of all, for every $v \in V$, $\sigma_v(x, y)$ has to be monotonically decreasing in the number of resources x that are placed at node v and monotonically increasing in the number of demands y at v.

$$\forall v \in V\, \forall x, y \in \mathbb{N}_0 \quad : \quad \sigma_v(x, y) \geq \sigma_v(x+1, y) \tag{7}$$

$$\forall v \in V\, \forall x, y \in \mathbb{N}_0 \quad : \quad \sigma_v(x, y) \leq \sigma_v(x, y+1) \tag{8}$$

Further, the effect of adding additional resources to a node v should become smaller with the number of resources and it should not decrease if the number of demands gets larger. Therefore, for all $v \in V$ and all $x, y \in \mathbb{N}_0$, we have

$$\sigma_v(x, y) - \sigma_v(x+1, y) \quad \geq \quad \sigma_v(x+1, y) - \sigma_v(x+2, y) \tag{9}$$

$$\sigma_v(x, y) - \sigma_v(x+1, y) \quad \leq \quad \sigma_v(x, y+1) - \sigma_v(x+1, y+1) \tag{10}$$

Finally, we define two constants characterizing the properties of the service cost function. We first define a constant $\Delta_{\max} > 0$, denoting the maxmum effect a single demand can have on the service cost of a node,

$$\Delta_{\max} := \max_{v \in V, f, r \in \mathbb{N}_0} \{\sigma_v(f, r+1) - \sigma_v(f, r)\} \overset{(10)}{=} \max_{v \in V, r \in \mathbb{N}_0} \{\sigma_v(0, r+1) - \sigma_v(0, r)\}.$$

Even though we do not explicitly bound $\Delta_{\max}$, we think of $\Delta_{\max}$ as being bounded by some constant. In order to capture that also an optimal offline algorithm has to pay at least some minimal cost when moving any resources, we define

$$\sigma_{\min} := \min_{F \text{ feasible}, F \neq F_0} S_0(F) = \min_{F = \{(v, f_v)\} \text{ feasible}, F \neq F_0} \sum_{v \in V} \sigma_v(f_v, 0)$$

as the minimum cost any algorithm has to pay for a feasible solution $F \neq F_0$ (i.e., if at least one resource has been moved away from its initial node).

Mathematical Notation: We conclude the section by defining some notation we use throughout the manuscript. For two integers $a \leq b$, $[a, b]$ denotes the set of all integers between a and b (including a and b). That is, $[a, b] := \{i \in \mathbb{Z} : a \leq i \leq b\}$. Further, for an integer $a \geq 1$, we use $[a]$ as a short form to denote $[a] := [1, a]$. Further in the following, whenever clear from the context, we omit the superscript $\mathcal{A}$in the algorithm-dependent quantities defined above.

3 Contributions

In Section 4.1, we describe a simple, deterministic online algorithm $\mathcal{A}$ with the following properties. For two parameters $\alpha \geq 1$ and $\beta \geq 0$, $\mathcal{A}$ guarantees that at all times $t \geq 0$, $S_t^{\mathcal{A}} < \alpha S_t^* + \beta$. The algorithm $\mathcal{A}$ achieves this while keeping the overall movement cost small. More precisely, we prove the following main theorem.

Theorem 3.1 (Upper Bound)**.** *For any $\alpha \geq 1$ and $\beta \geq 0$ satisfying* (12)*, there is a deterministic online algorithm $\mathcal{A}$, such that for all times $t \geq 0$, we have $S_t^{\mathcal{A}} < \alpha S_t^* + \beta$ and where the total movement cost $M_t^{\mathcal{A}}$ is bounded as follows.*

- *If $\alpha = 1$, for any $\ell \geq 1$, $\varepsilon > 0$, and $\beta \geq k(3k)^{1/\ell}/\varepsilon$, we have*

$$M_t^{\mathcal{A}} \leq \varepsilon \cdot S_t^* + O(\ell k).$$

- *For $\alpha \geq 1 + \varepsilon$, for some constant $\varepsilon > 0$, and any $\beta \geq 0$, we have*

$$M_t^{\mathcal{A}} \leq k \cdot O\left(1 + \log_\alpha S_t^* + \min\left\{\frac{\log k}{\log\log k}, \log_\alpha k\right\} + \log_\alpha \frac{k}{\sigma_{\min} + \beta}\right).$$

Note that by choosing $\alpha > 1$, the dependency of the movement cost $M_t^{\mathcal{A}}$ on the optimal service cost S_t^* is only logarithmic. As a bound on the total cost, the theorem implies the following corollary.

Corollary 3.2. *For any constant $0 < \delta \leq 1$ and every $\varepsilon \geq 1/\log^{1-\delta} k$, there is a deterministic algorithm $\mathcal{A}$ satisfying*

$$S_t^{\mathcal{A}} + M_t^{\mathcal{A}} \leq (1 + \varepsilon) S_t^* + O\left(\frac{k}{\delta} \cdot \frac{\log k}{\log\log k}\right).$$

Proof. Follows from the first bound of Theorem 3.1 by choosing $\ell = \Theta\left(\frac{1}{\delta} \cdot \frac{\log k}{\log\log k}\right)$. □

We also prove an almost matching lower bound. Our lower bound holds for any deterministic online algorithm and even for a particularly simple (and natural) scenario, where the service cost at a node with at least 1 resource is 0 and where serving y demands at a node without resource costs y. When modeling the service cost like this, we need to explicitly keep track of the movement cost. We use M_t^* to denote $\chi(F_0, F_t^*)$, the optimal movement cost for the first t demands. For a description of how to incorporate the movement cost of an offline algorithm into the general cost function, see Appendix A. As a consequence our lower bound also applies to the online version of the movement minimization problem in standard metric mobile facility location, even for uniform metrics [1, 7].

Theorem 3.3 (Lower Bound)**.** *Let $\alpha \geq 1$ and $\beta \geq 0$ be two parameters. Assume that there exists an online algorithm $\mathcal{A}$ for the above restricted setting and assume that $\mathcal{A}$ at all time guarantees that $S_t^{\mathcal{A}} < \alpha(S_t^* + M_t^*) + \beta$. Then, there exists an execution and a time $t \geq 0$ such that the total movement cost $M_t^{\mathcal{A}}$ of $\mathcal{A}$ can be lower bounded as follows.*

- *If $\alpha = 1$, for any $\ell \geq 1$, $\varepsilon > 0$, and $\beta \leq k(k/3)^{1/\ell}/\varepsilon$, we have*
$$M_t^{\mathcal{A}} \geq \varepsilon \cdot S_t^* + \Omega(\ell k).$$
- *For $\alpha \geq 1 + \varepsilon$, for some constant $\varepsilon > 0$, and any $\beta \geq 0$, we have*
$$M_t^{\mathcal{A}} \geq k \cdot \Omega(1 + \log_\alpha S_t^*).$$

As a corollary, we obtain:

Corollary 3.4. *Let $0 < \delta \leq 1$ and $\varepsilon > 1/\log^{1-\delta} k$ and assume that there exists some deterministic online algorithm $\mathcal{A}$ that guarantees that for all $t \geq 0$, $S_t^{\mathcal{A}} + M_t^{\mathcal{A}} \leq (1+\varepsilon)(S_t^* + M_t^*) + \beta$ for some $\beta \geq 0$. Then, we have $\beta = \Omega\left(\frac{k}{\delta} \cdot \frac{\log k}{\log\log k}\right)$.*

4 Online Algorithm

4.1 Algorithm Description

The goal of our algorithm is two-fold. On the one hand, we want to guarantee that the service cost of the algorithm is always within some fixed bounds of the optimal service cost. On the other hand, we want to achieve this while keeping the overall movement cost low. Specifically, for two parameters α and β, we guarantee that at all times
$$S_t < \alpha \cdot S_t^* + \beta, \tag{11}$$
where S_t denotes the total service cost of the algorithm at time t. Condition (11) is maintained in the most straightforward greedy manner. Whenever after a new demand arrives, (11) is not satisfied, the algorithm greedily move resources until (11) holds again. Hence, as long as (11) does not hold, the algorithm moves a resource that reduces the total service cost as much as possible. The algorithm stops moving any resources as soon as the validity of (11) is restored. For our analysis to work, we need to guarantee that the algorithm tolerates a deviation from the optimal service cost that is at least in the order of the effect of adding a single demand. We therefore assume that the parameters α and β are chosen such that
$$\max\{(\alpha - 1)\sigma_{\min}, \beta\} \geq \Delta_{\max}. \tag{12}$$

Whenever the algorithm moves a resource, it does a best possible move, i.e., a move that achieves the best possible service cost improvement. Thus, the algorithm always moves a resource from a node where

removing a resource is as cheap as possible to a node where adding a resource reduces the cost as much as possible. Therefore, for each movement m, we have

$$v_m^{src} \in \arg\min_{v \in V} \{\sigma_{v,\tau_m}(f_{v,m-1} - 1) - \sigma_{v,\tau_m}(f_{v,m-1})\} \text{ and} \quad (13)$$

$$v_m^{dst} \in \arg\max_{v \in V} \{\sigma_{v,\tau_m}(f_{v,m-1}) - \sigma_{v,\tau_m}(f_{v,m-1} + 1)\}, \quad (14)$$

where $\arg\min_v$ and $\arg\max_v$ denote the sets of nodes minimizing and maximizing the respective terms.

4.2 Analysis Overview

While the algorithm itself is quite simple, its analysis turns out relatively technical. We thus first describe the key steps of the analysis by discussing a simple case. We assume that the service cost at any node is equal to 0 if there is at least one resource at the node and the service cost is equal to the number of demands at the node otherwise. Further, we assume that we run the algorithm of 4.1 with parameters $\alpha = 1$ and $\beta = 0$, i.e. after each demand arrives, the algorithm moves to a solution with optimal service cost. Note that these parameter settings violate Condition (12) and we will therefore get a weaker bound than the one promised by Theorem 3.1.

First, note that in the described simple scenario, the algorithm clearly never puts more than one resource to the same node. Further, whenever the algorithm moves a resource from a node u to a node v, the overall service cost has to strictly decrease and thus, the number of demands at node v is larger than the number of demands at node u. Consider some point in time t and let $d_{\min}(t) := \min_{v \in V: f_{v,t}=1} d_{v,t}$ be the minimum number of demands among the nodes v with a resource at time t. Hence, whenever at a time t, the algorithm moves a resource from a node u to a node v, node u has at least $d_{\min}(t)$ demands and consequently, node v has at least $d_{\min(t)} + 1$ demands. Further, if at some later time $t' > t$, the resource at node v is moved to some other node w, because the algorithm always removes a resource from a node with as few demands as possible, we have $d_{\min}(t') \geq d_{\min}(t) + 1$. Consequently, if in some time interval $[t_1, t_2]$, there is some resource that is moved more than once, we know that $d_{\min}(t_1) < d_{\min}(t_2)$. In our analysis, we use this observation to partition time into phases, where the first phase starts at time 0 and where phases are maximal time intervals in which each resource is moved at most once (cf. Def. 4.1 in the formal analysis of the algorithm). Note that clearly, as there are only k different resources, the number of movements in a phase can be at most k (cf. Observation 4.2).

The above argument implies that after each phase $d_{\min}$ increases by at least one and therefore at any time t in phase p, we have $d_{\min}(t) \geq p - 1$ and at the end of phase p, we have $d_{\min}(t) \geq p$. In Section 4.3, the more general form of this statement appears in Lemma 4.3. There, γ_p is defined to be the smallest service cost improvement of any movement in phase p ($\gamma_p = 1$ in the simple case considered here), and Lemma 4.3 shows that $d_{\min}$ grows by at least γ_p in phase p.

Assume that at some time t in phase p, a resource is moved from a node u to a node v. Because node u already had its resource at the end of phase $p-1$, we have $d_{u,t} = d_{\min}(t) \geq p-1$. Consequently, at the end of phase p, there is at least one node (the source of the last movement) that has no resource and at least $p-1$ demands. The corresponding (more technical) statement in our general analysis appears in Lemma 4.5.

Our algorithm guarantees that at all times, the service cost is within fixed bounds of the optimal service cost (in the simple case here, the service cost is always equal to the optimal service cost). Knowing that there are nodes with many demands and no resources, therefore allows to lower bound the optimal service cost. In the general case, this is done by Lemmas 4.8 and 4.9. In the simple case, considered here, as at the end of phase p, there are k nodes with at least p demands (the nodes that have resources) and there is at least one additional node with at least $p-1$ demands, we know that at the end of phase p, the optimal service cost is at least $p-1$. Consequently, since the number of movements per phase is at most k, the total number of movements of the algorithm is at most k times the cost of an optimal offline algorithm.

4.3 Algorithm Analysis

We are now going to analyze the algorithm of Section 4.1. In the following, whenever we refer to the algorithm introduced Section 4.1, we omit the superscript $\mathcal{A}$. For the analysis of the described algorithm, we partition the movements into phases $p = 1, 2, \ldots$, where roughly speaking, a phase is a maximal consecutive sequence of movements in which no resource is moved twice. We use m_p to denote the first movement of phase p (for $p \in \mathbb{N}$). Formally, the phases are defined as follows.

Definition 4.1 (Phases)**.** *The movements are divided into phases $p = 1, 2, \ldots$, where phase p starts with movement m_p and ends with movement $m_{p+1} - 1$. We have $m_1 = 1$, i.e., the first phase starts with the first movement. Further for every $p > 1$, we define*

$$m_p := \min\left\{ m > m_{p-1} \;:\; \exists\, m' \in [m_{p-1}, m-1] \text{ s.t. } v_m^{src} = v_{m'}^{dst} \right\} - 1. \tag{15}$$

For a phase $p \geq 1$, let $\lambda_p := m_{p+1} - m_p$ be the number of movements of Phase p. As an immediate consequence of the above definition, we obtain that the maximum number of movements in each phase is at most k.

Observation 4.2. *For each Phase $p \geq 1$, we have $\lambda_p \leq k$.*

Proof. Let $m > m_p$ and consider the movements $[m_p, m]$. We prove that if $m < m_{p+1}$, no two the movements in $[m_p, m]$ move the same resource. The claim then follows because there are only k resources. For the sake of contradiction, assume that there is some resource i that is moved more than once and let m' and m'' ($m', m'' \in [m_p, m]$, $m' < m''$) be the first two movements in $[m_p, m]$, where resource i is moved. We clearly have $v_{m'}^{dst} = v_{m''}^{src}$ and Def. 4.1 thus leads to a contradiction to the assumption that $m < m_{p+1}$. □

The algorithm moves resources in order to improve the service cost. For a given movement m, we use $\gamma(m) > 0$ to denote service cost *improvement* of m:

$$\begin{aligned} \gamma(m) &:= S_{\tau_m}(F_{m-1}) - S_{\tau_m}(F_m) \\ &= \left(\sigma_{v_m^{dst},\tau_m}(f_{v_m^{dst},m-1}) - \sigma_{v_m^{dst},\tau_m}(f_{v_m^{dst},m})\right) \\ &\quad - \left(\sigma_{v_m^{src},\tau_m}(f_{v_m^{src},m}) - \sigma_{v_m^{src},\tau_m}(f_{v_m^{src},m-1})\right). \end{aligned} \tag{16}$$

For each Phase p, we define the improvement γ_p of p and the cumulative improvement Γ_p by Phase p as follows

$$\gamma_p := \min_{m \in [m_p, m_{p+1}-1]} \gamma(m) \quad \text{and} \quad \Gamma_p := \sum_{i=1}^{p} \gamma_i, \quad \Gamma_0 := 0,\ \gamma_0 := 0. \tag{17}$$

We are now ready to prove our first technical lemma, which lower bounds the cost of removing resources from nodes with resources at any point in the execution.

Lemma 4.3. *Let m be a movement and, $F = \{(v, f_v) : v \in V\}$ be the solution of the algorithm at any point in the execution after movement m and let $t \geq \tau_m$ be the time at which the solution F occurs. Then, for all times $t' \geq t$ and for all nodes $v \in V$, if $f_v > 0$ it holds that $\sigma_{v,t'}(f_v - 1) - \sigma_{v,t'}(f_v) \geq \Gamma_{p-1}$, where p is the phase in which movement m occurs.*

Proof. We will show that for each resource movement $m \in \mathbb{N}$ of the algorithm, it holds that

$$\forall v \in V \;:\; f_{v,m} > 0 \implies \sigma_{v,\tau_m}(f_{v,m} - 1) - \sigma_{v,\tau_m}(f_{v,m}) \geq \Gamma_{p-1}, \tag{18}$$

where p is the phase in which movement m occurs (i.e., the claim of the lemma holds immediately after movement m). The lemma then follows because a) any solution $\{(v, f_v) : v \in V\}$ occurring after movement m is the solution $F_{m'}$ for some movement $m' \geq m$, b) the values Γ_{p-1} are monotonically increasing with p, and c) by (10), for all $v \in V$, the value $\sigma_{v,t}(f-1) - \sigma_{v,t}(f)$ is monotonically non-decreasing with t.

It therefore remains to prove (18) for every m, where p is the phase of movement m. We prove a slightly stronger statement. Generally, for a movement m' and a phase p', let $V^{dst}_{p',m'}$ be the set of nodes that have received a new resource by some movement $m'' \leq m'$ of Phase p'. Hence,

$$V^{dst}_{p',m'} := \left\{ v \in V : \exists \text{ movement } m'' \leq m' \text{ of Phase } p' \text{ s.t. } v^{dst}_{m''} = v \right\}.$$

We show that in addition to (18), it also holds that

$$\forall v \in V^{dst}_{p,m} \ : \ f_{v,m} > 0 \implies \sigma_{v,\tau_m}(f_{v,m}-1) - \sigma_{v,\tau_m}(f_{v,m}) \geq \Gamma_p. \tag{19}$$

We prove (18) and (19) together by using induction on m.

Induction Base ($m = 1$): The first movement occurs in Phase 1. By (17), $\Gamma_0 = 0$ and by (7), we also have $\sigma_{v,t}(f-1) - \sigma_{v,t}(f) \geq 0$ for all times $t \geq 0$, all nodes $v \in V$, and all $f \geq 1$. Inequality (18) therefore clearly holds for $m = 1$. It remains to show that also (19) holds for $m = 1$. We have $V^{dst}_{1,1} = \left\{v^{dst}_1\right\}$ and showing (19) for $m = 1$ therefore reduces to showing that $\sigma_{v^{dst}_1,\tau_1}(f_{v^{dst}_1,1}-1) - \sigma_{v^{dst}_1,\tau_1}(f_{v^{dst}_1,1}) \geq \Gamma_1 = \gamma_1$, which follows directly from (16) and (17).

Induction Step ($m > 1$): We first show that Inequalities (18) and (19) hold immediately before movement m and thus,

$$\forall v \in V \ : \ f_{v,m-1} > 0 \ \Rightarrow \ \sigma_{v,\tau_m}(f_{v,m-1}-1) - \sigma_{v,\tau_m}(f_{v,m-1}) \geq \Gamma_{p-1}, \tag{20}$$

$$\forall v \in V^{dst}_{p,m-1} \ : \ f_{v,m-1} > 0 \ \Rightarrow \ \sigma_{v,\tau_m}(f_{v,m-1}-1) - \sigma_{v,\tau_m}(f_{v,m-1}) \geq \Gamma_{p}. \tag{21}$$

If m is not the first movement of Phase p, Inequalities (20) and (21) follow directly from the induction hypothesis (for $m-1$) and from (10). Let us therefore assume that m is the first movement of Phase p. Note that in this case $V^{dst}_{p,m-1} = \emptyset$ and (21) therefore trivially holds. Because $m > 1$, we know that in this case $p \geq 2$. From the induction hypothesis and from (10), we can therefore conclude that for every node $v \in V^{dst}_{p-1,m-1}$ (every node v that is the destination of some resource movement in Phase $p-1$), we have $\sigma_{v,\tau_m}(f_{v,m-1}-1) - \sigma_{v,\tau_m}(f_{v,m-1}) \geq \Gamma_{p-1}$. Note that for all these nodes, we have $f_{v,m-1} > 0$. Because m is the first movement of Phase p, Definition 4.1 implies that $v^{src}_m \in V^{dst}_{p-1,m-1}$. Applying (13), we get that for all $v \in V$, $\sigma_{v,\tau_m}(f_{v,m-1}-1) - \sigma_{v,\tau_m}(f_{v,m-1}) \geq \sigma_{v^{src}_m,\tau_m}(f_{v^{src}_m,m-1}-1) - \sigma_{v^{src}_m,\tau_m}(f_{v^{src}_m,m-1}) \geq \Gamma_{p-1}$ and therefore (20) also holds if $m \geq 2$ is the first movement of some phase.

We can now prove (18) and (19). For all nodes $v \notin \left\{v^{src}_m, v^{dst}_m\right\}$, we have $f_{v,m} = f_{v,m-1}$ and we further have $V^{dst}_{p,m} = V^{dst}_{p,m-1} \cup \left\{v^{dst}_m\right\}$. For $v \notin \left\{v^{src}_m, v^{dst}_m\right\}$, (18) and (19) therefore directly follow from (20) and (21), respectively. For the two nodes involved in movement m, first note that $v^{src}_m \notin V^{dst}_{p,m-1}$. It therefore suffices to show that

$$f_{v^{src}_m,m} = 0 \quad \text{or} \quad \sigma_{v^{src}_m,\tau_m}(f_{v^{src}_m,m}-1) - \sigma_{v^{src}_m,\tau_m}(f_{v^{src}_m,m}) \ \geq \ \Gamma_{p-1}, \tag{22}$$

$$\text{as well as} \quad \sigma_{v^{dst}_m,\tau_m}(f_{v^{dst}_m,m}-1) - \sigma_{v^{dst}_m,\tau_m}(f_{v^{dst}_m,m}) \ \geq \ \Gamma_{p}. \tag{23}$$

We have $f_{v^{src}_m,m} = f_{v^{src}_m,m-1} - 1$ and $f_{v^{dst}_m,m} = f_{v^{dst}_m,m-1} + 1$. Inequality (22) therefore directly follows from (20) and from (9). For (23), we have

$$\begin{aligned}
&\sigma_{v^{dst}_m,\tau_m}(f_{v^{dst}_m,m}-1) - \sigma_{v^{dst}_m,\tau_m}(f_{v^{dst}_m,m}) \\
&\quad \overset{(16)}{=} \ \sigma_{v^{src}_m,\tau_m}(f_{v^{src}_m,m}) - \sigma_{v^{src}_m,\tau_m}(f_{v^{src}_m,m}-1) + \gamma(m) \\
&\quad \overset{(20)}{\geq} \ \Gamma_{p-1} + \gamma(m) \ \overset{(17)}{\geq} \ \Gamma_p.
\end{aligned}$$

This completes the proof of (18) and (19) and thus the proof of the lemma. □

For each phase number p, let $\theta_p := \tau_{m_p}$ be the time of the the first movement m_p of Phase p. Before continuing, we give lower and upper bounds on γ_p, the improvement of Phase p. For all $p \geq 1$, we define

$$\eta_p := (\alpha - 1) \cdot S^*_{\theta_p} + \beta. \tag{24}$$

Lemma 4.4. *Let m be a movement of Phase p and let $F^* = \{(v, f^*_v) : v \in V\} \in \mathcal{F}^*(\tau_m)$ be an optimal feasible solution at time τ_m. We then have*

$$\frac{\eta_p}{\chi(F_{m-1}, F^*)} \leq \gamma(m) < \eta_{p+1} + \Delta_{\max}.$$

Proof. For the upper bound, observe that we have

$$\gamma(m) \leq S_{\tau_m}(F_{m-1}) - S^*_{\tau_m}$$

as clearly the service cost cannot be improved by a larger amount. Because at all times t, the algorithm keeps the service cost below $\alpha S^*_t + \beta$, we have $S_{\tau_m - 1}(F_{m-1}) < \alpha S^*_{\tau_m - 1} + \beta \leq \alpha S^*_{\tau_m} + \beta$. As the single demand arriving at time τ_m can only increase the service cost by at most $\Delta_{\max}$, the upper bound on $\gamma(m)$ follows from (24) and because $S^*_{\tau_m} \leq S^*_{\theta_{p+1}}$.

For the lower bound on $\gamma(m)$, we need to prove that $\chi(F_{m-1}, F^*) \geq \eta_p / \gamma(m)$. Because the algorithm moves a resource at time τ_m, we know that $S_{\tau_m}(F_{m-1}) \geq \alpha S_{\tau_m}(F^*) + \beta$ and applying the Definition (24) of η_p, we thus have $S_{\tau_m}(F_{m-1}) - S_{\tau_m}(F^*) \geq \eta_p$. Intuitively, we have $\chi(F_{m-1}, F^*) \geq \eta_p / \gamma(m)$ because the algorithm always chooses the best possible movement and thus every possible movement improves the overall service cost by at most $\gamma(m)$. Thus, the number of movements needs to get from F_{m-1} to an optimal solution F^* has to be at least $\eta_p / \gamma(m)$. For a formal argument, assume that we are given a sequence of $\ell := \chi(F_{m-1}, F^*)$ movements that transform solution F_{m-1} into solution F^*. For $i \in [\ell]$, assume that the i^{th} of these movements moves a resource from node u_i to node v_i. Further, for any $i \in [\ell]$ let f_i be the number of resources at node u_i and let f'_i be the number of resources at node v_i before the i^{th} of these movements. Because the sequence of movements is minimal to get from F_{m-1} to F^*, we certainly have $f_i \leq f_{u_i, m-1}$ and $f'_i \geq f_{v_i, m-1}$. For the service cost improvement γ of the i^{th} of these movements, we therefore obtain

$$\begin{aligned}
\gamma &= \big(\sigma_{v_i,\tau_m}(f'_i) - \sigma_{v_i,\tau_m}(f'_i + 1)\big) - \big(\sigma_{u_i,\tau_m}(f_i - 1) - \sigma_{u_i,\tau_m}(f_i)\big) \\
&\overset{(9)}{\leq} \big(\sigma_{v_i,\tau_m}(f_{v_i,m-1}) - \sigma_{v_i,\tau_m}(f_{v_i,m-1} + 1)\big) \\
&\qquad - \big(\sigma_{u_i,\tau_m}(f_{u_i,m-1} - 1) - \sigma_{u_i,\tau_m}(f_{u_i,m-1})\big) \\
&\leq \gamma(m).
\end{aligned}$$

The last inequality follows from (13),(14), and (16). As the sum of the ℓ service cost improvements has to be at least η_p, we obtain $\ell = \chi(F_{m-1}, F^*) \geq \eta_p / \gamma(m)$ as claimed. □

We can now lower bound the distribution of demands at the time of each movement.

Lemma 4.5. *Let m be a movement of Phase p (for $p \geq 1$). Then, there are integers $\psi_v \geq 0$ for all nodes $v \in V$ such that*

$$\sum_{v \in V} \psi_v \geq k + \frac{\eta_p}{\gamma(m)} \quad \textit{and}$$

$$\forall t \geq \tau_m \; \forall v \in V : \psi_v > 0 \Longrightarrow \sigma_{v,t}(\psi_v - 1) - \sigma_{v,t}(\psi_v) \geq \Gamma_{p-1}.$$

Proof. It suffices to prove the statement for $t = \tau_m$. For larger t, the claim then follows from (10). Consider an optimal feasible solution

$$F^* = \{(v, f_v^*) : v \in V\} \in \mathcal{F}^*(\tau_m)$$

at the time τ_m of movement m. Let us further consider the solution F_{m-1} of the algorithm immediately before movement m. Consider a pair of nodes u and v such that $f_u^* > f_{u,m-1}$ and $f_{v,m-1} > f_v^*$. By the optimality of F^*, we have

$$\sigma_{u,\tau_m}(f_u^* - 1) - \sigma_{u,\tau_m}(f_u^*) \geq \sigma_{v,\tau_m}(f_{v,m-1} - 1) - \sigma_{v,\tau_m}(f_{v,m-1}). \tag{25}$$

Otherwise, moving a resource from u to v would (strictly) improve the solution F^*. By Lemma 4.3, we have $\sigma_{v,\tau_m}(f_{v,m-1} - 1) - \sigma_{v,\tau_m}(f_{v,m-1}) \geq \Gamma_{p-1}$ for all nodes v for which $f_{v,m-1} > 0$. Together with (25), for all $v \in V$ for which $\max\{f_{v,m-1}, f_v^*\} > 1$, we obtain

$$\sigma_{v,\tau_m}(\max\{f_{v,m-1}, f_v^*\} - 1) - \sigma_{v,\tau_m}(\max\{f_{v,m-1}, f_v^*\}) \geq \Gamma_{p-1}. \tag{26}$$

To prove the lemma, it therefore suffices to show that $\sum_{v \in V} \max\{f_{v,m-1}, f_v^*\} \geq k + \eta_p/\gamma(m)$, as we can then set $\psi_v := \max\{f_{v,m-1}, f_v^*\}$ and (26) implies the claim of the lemma. By (4), we have

$$\sum_{v \in V} \max\{f_{v,m-1}, f_v^*\} = k + \sum_{v \in V} \max\{0, f_v^* - f_{v,m-1}\} = k + \chi(F_{m-1}, F^*).$$

We therefore need that $\chi(F_{m-1}, F^*) \geq \eta_p/\gamma(m)$, which follows from Lemma 4.4. □

In the next lemma, we derive a lower bound on $S_{\theta_p}^*$, the service cost of an optimal offline algorithm when Phase p starts. For each phase $p \geq 1$, we first define $\overline{S}_p$ as follows.

$$\text{For } p \geq 3 : \overline{S}_p := \left(1 + (\alpha - 1)\frac{\gamma_{p-2}}{\gamma_{p-1}}\right) \cdot \overline{S}_{p-1} + \frac{\gamma_{p-2}}{\gamma_{p-1}}\beta, \text{ and } \overline{S}_1 := \overline{S}_2 := \sigma_{\min}. \tag{27}$$

Lemma 4.6. *For all $p \geq 1$, we have $S_{\theta_p}^* \geq \overline{S}_p$.*

Proof. We prove the lemma by induction on p.

Induction Base ($p = 1, 2$): Considering the definition of $\sigma_{\min}$ in Section 2 we have $S_{\theta_1}^* \geq \sigma_{\min}$ and since $\overline{S}_1 = \overline{S}_2 = \sigma_{\min}$, we get $S_{\theta_2}^* \geq S_{\theta_1}^* \geq \overline{S}_2 = \overline{S}_1$.

Induction Step ($p > 2$): We use the induction hypothesis to assume that the claim of the lemma is true up to Phase p and we prove that it also holds for Phase $p + 1$. Therefore by the induction hypothesis, for all $i \in [p]$,

$$S_{\theta_p}^* \geq \overline{S}_p. \tag{28}$$

For all $i \in [p]$, we define $\overline{\eta}_i := (\alpha - 1)\overline{S}_i + \beta$ and $\delta_i := \max\left\{\frac{\overline{\eta}_{i+1}}{\gamma_{i+1}}, \cdots, \frac{\overline{\eta}_p}{\gamma_p}\right\}$ As a consequence of (24) and (28), we get that $\eta_i \geq \overline{\eta}_i$ for all $i \in [p]$. In the following, let $p' \in [2, p]$ be some phase. Lemma 4.5 implies that after the last movement m of Phase p', there are non-negative integers ψ_v (for $v \in V$) such that $\sum_{v \in V} \psi_v \geq k + \eta_{p'}/\gamma_{p'} \geq \overline{\eta}_{p'}/\gamma_{p'}$ and for all times $t \geq \tau_m$, for all $v \in V$ for which $\psi_v > 0$, $\sigma_{v,t}(\psi_v - 1) - \sigma_{v,t}(\psi_v) \geq \Gamma_{p'-1}$. As there are only k resources for any feasible solution $F = \{(v, f_v)\}$, we have $\sum_{v \in V} f_v = k$ and therefore $\sum_{v \in V}(\psi_v - f_v) \geq \overline{\eta}_{p'}/\gamma_{p'}$. For any $v \in V$ for which $\psi_v > f_v$, by using (9), we get $\sigma_{v,t}(f_v) \geq (\psi_v - f_v)\Gamma_{p'-1}$. Hence, after the last movement of Phase p', for any feasible solution

F, we have $S_t(F) \geq S_t^* \geq \frac{\overline{\eta}_{p'}}{\gamma_{p'}}\Gamma_{p'-1}$. At the beginning of Phase $p+1$ (for $p \geq 2$), the total service cost of an optimal offline algorithm therefore is

$$S^*_{\theta_{p+1}} \geq \max_{p' \in [2,p]} \frac{\overline{\eta}_{p'}}{\gamma_{p'}}\Gamma_{p'-1} \geq \delta_{p-1}\Gamma_{p-1} + \sum_{i=1}^{p-2}(\delta_i - \delta_{i+1}) \cdot \Gamma_i = \sum_{i=1}^{p-1}\gamma_i \cdot \delta_i. \tag{29}$$

We define ζ_i for all $i \in [3,p]$ as follows:

$$\zeta_i := \sum_{j=1}^{i-2}\gamma_j \cdot \delta_j. \tag{30}$$

Using the definition of δ_i, we thus have

$$\zeta_{p+1} = \zeta_p + \gamma_{p-1}\delta_{p-1} = \zeta_p + \overline{\eta}_p\frac{\gamma_{p-1}}{\gamma_p}.$$

Considering the definition of $\overline{\eta}_i$ we get

$$\zeta_{p+1} = \zeta_p \cdot \left(1 + (\alpha - 1)\frac{\gamma_{p-1}}{\gamma_p}\right) + \beta \cdot \frac{\gamma_{p-1}}{\gamma_p}.$$

We therefore have $\zeta_{p+1} = \overline{S}_{p+1}$ directly from (27) and thus the claim of the lemma follows. □

In order to explicitly lower bound the optimal service cost after p phases, we need the following technical statement.

Lemma 4.7. *Let $\ell \geq 2$ be an integer and consider a sequence $c_1, c_2, \ldots, c_\ell > 0$ of ℓ positive real numbers and let $c_{\max} = \max_{i\in[\ell]} c_i$ and $c_{\min} = \min_{i\in[\ell]} c_i$. Further, let $\lambda \geq 0$ be an arbitrary non-negative real number. We have*

$$\text{(I)} \quad \sum_{i=2}^{\ell}\frac{c_{i-1}}{c_i} \geq (\ell-1)\cdot\left(\frac{c_{\min}}{c_{\max}}\right)^{\frac{1}{\ell-1}},$$

$$\text{(II)} \quad \prod_{i=2}^{\ell}\left(1+\lambda\frac{c_{i-1}}{c_i}\right) \geq \left(1+\lambda\left(\frac{c_{\min}}{c_{\max}}\right)^{\frac{1}{\ell-1}}\right)^{\ell-1}.$$

Proof. The first part of the claim follows from the means inequality (the fact that the arithmetic mean is larger than or equal to the geometric mean). In the following, we nevertheless directly prove both parts together. We let $\boldsymbol{x} = (x_1, \ldots, x_\ell) \in \mathbb{R}^\ell$ be a vector ℓ real variables and we define multivariate functions $f(\boldsymbol{x}) : \mathbb{R}^\ell \to \mathbb{R}$ and $g(\boldsymbol{x}) : \mathbb{R}^\ell \to \mathbb{R}$ as follows:

$$f(\boldsymbol{x}) := \sum_{i=2}^{\ell}\frac{x_{i-1}}{x_i} \quad \text{and} \quad g(\boldsymbol{x}) := \prod_{i=2}^{\ell}\left(1+\lambda\frac{x_{i-1}}{x_i}\right).$$

We further define $X \subset \mathbb{R}^\ell$ as $X := \left\{(z_1, \ldots, z_\ell) \in \mathbb{R}^\ell \,|\, \forall i \in [\ell] : c_{\min} \leq z_i \leq c_{\max}\right\}$. We need to show that for $\boldsymbol{x} \in X$, $f(\boldsymbol{x})$ and $g(\boldsymbol{x})$ are lower bounded by the right-hand sides of Inequalities (I) and (II) above, respectively. Note that X is a closed subset of $\mathbb{R}^\ell$ and because $c_{\min} > 0$, both functions $f(\boldsymbol{x})$ and $g(\boldsymbol{x})$ are

continuous when defined on X. The minimum for $\boldsymbol{x} \in X$ is therefore well-defined for both $f(\boldsymbol{x})$ and $g(\boldsymbol{x})$. We show that both $f(\boldsymbol{x})$ and $g(\boldsymbol{x})$ attain their minimum for

$$\boldsymbol{x^*} := (x_1^*, \ldots, x_\ell^*), \quad \text{where } \forall i \in [\ell] : x_i^* = c_{\min} \cdot \left(\frac{c_{\max}}{c_{\min}}\right)^{\frac{i-1}{\ell-1}}.$$

Note that $\boldsymbol{x^*}$ is the unique solution $\boldsymbol{x} \in X$ to the following system of equations

$$x_1 = c_{\min}, \quad x_\ell = c_{\max}, \quad \forall i \in \{2, \ldots, \ell-1\} : x_i \in \frac{x_{i-1}}{x_i} = \frac{x_i}{x_{i+1}}. \tag{31}$$

Because we know that $\min_{\boldsymbol{x} \in X} f(\boldsymbol{x}) = f(\boldsymbol{x^*})$ and $\min_{\boldsymbol{x} \in X} g(\boldsymbol{x}) = g(\boldsymbol{x^*})$, it is therefore sufficient to show that for any $\boldsymbol{y} \in X$ that does not satisfy (31), $f(\boldsymbol{y})$ and $g(\boldsymbol{y})$ are not minimal. Let us therefore consider a vector $\boldsymbol{y} = (y_1, \ldots, y_\ell) \in X$ that does not satisfy (31). First note that both $f(\boldsymbol{x})$ and $g(\boldsymbol{x})$ are strictly monotonically increasing in x_1 and strictly monotonically decreasing in x_ℓ. If either $y_1 > c_{\min}$ or $y_\ell < c_{\max}$, it is therefore clear that $f(\boldsymbol{y})$ and $g(\boldsymbol{y})$ are both not minimal (over X). Let us therefore assume that $y_1 = c_{\min}$ and $y_\ell = c_{\max}$. From the assumption that $\boldsymbol{y}$ does not satisfy (31), we then have an $i_0 \in \{2, \ldots, \ell-1\}$ for which $\frac{y_{i_0-1}}{y_{i_0}} \neq \frac{y_{i_0}}{y_{i_0+1}}$ and thus $y_{i_0} \neq \sqrt{y_{i_0-1} y_{i_0+1}}$. We define a new vector $\boldsymbol{y'} = (y'_1, \ldots, y'_\ell) \in X$ as follows. We have $y'_{i_0} = \sqrt{y_{i_0-1} y_{i_0+1}}$ and $y'_i = y_i$ for all $i \neq i_0$ and we will show that $f(\boldsymbol{y'}) < f(\boldsymbol{y})$ and $g(\boldsymbol{y'}) < g(\boldsymbol{y})$. Define

$$C := \prod_{i \in [2,\ell] \setminus \{i_0, i_0+1\}} \left(1 + \lambda \frac{y_{i-1}}{y_i}\right).$$

We then have

$$\begin{aligned}
f(\boldsymbol{y}) - f(\boldsymbol{y'}) &= \left(\frac{y_{i_0-1}}{y_{i_0}} + \frac{y_{i_0}}{y_{i_0+1}}\right) - \left(\frac{y_{i_0-1}}{y'_{i_0}} + \frac{y'_{i_0}}{y_{i_0+1}}\right) \\
g(\boldsymbol{y}) - g(\boldsymbol{y'}) &= \left[\left(1 + \lambda \frac{y_{i_0-1}}{y_{i_0}}\right) \cdot \left(1 + \lambda \frac{y_{i_0}}{y_{i_0+1}}\right) - \left(1 + \lambda \frac{y_{i_0-1}}{y'_{i_0}}\right) \cdot \left(1 + \lambda \frac{y'_{i_0}}{y_{i_0+1}}\right)\right] \cdot C \\
&= \left[\left(\frac{y_{i_0-1}}{y_{i_0}} + \frac{y_{i_0}}{y_{i_0+1}}\right) - \left(\frac{y_{i_0-1}}{y'_{i_0}} + \frac{y'_{i_0}}{y_{i_0+1}}\right)\right] \cdot \lambda C.
\end{aligned}$$

Note that $\lambda \geq 0$ and $C > 0$. In both cases, we therefore need to show that

$$\forall y_{i_0} \in [c_{\min}, c_{\max}] \setminus \left\{\sqrt{y_{i_0-1} y_{i_0+1}}\right\} : \left(\frac{y_{i_0-1}}{y_{i_0}} + \frac{y_{i_0}}{y_{i_0+1}}\right) > \left(\frac{y_{i_0-1}}{y'_{i_0}} + \frac{y'_{i_0}}{y_{i_0+1}}\right). \tag{32}$$

This follows because the function $h : [c_{\min}, c_{\max}] \to \mathbb{R}$, $h(z) := \frac{y_{i_0-1}}{z} + \frac{z}{y_{i_0+1}}$ is strictly convex for $z \in [c_{\min}, c_{\max}]$ and it has a stationary point at $z = \sqrt{y_{i_0-1} y_{i_0+1}} \in [c_{\min}, c_{\max}]$. □

As long as $(\alpha - 1) S^*_{\theta_p} < \beta$, the effect of the $(\alpha - 1) S^*_{\theta_p}$-term on η_p (and thus of the αS^*_t term in (11)) is relatively small. Let us therefore first analyze how the service cost grows by just considering terms that depends on β (and not on α). Note that the following lemma also holds if $\sigma_{\min} = 0$.

Lemma 4.8. *For all $p \geq 3$, we have*

$$S^*_{\theta_p} \geq \min\left\{\frac{\beta}{\alpha - 1},\ \beta \cdot (p-2) \cdot (3k)^{-\frac{1}{p-2}}\right\}.$$

Proof. Assume that $S^*_{\theta_p} < \beta/(\alpha-1)$ as otherwise the claim of the lemma is trivially true. By Lemma 4.6, using $\alpha \geq 1$, for all $p \geq 3$, we get $\overline{S}_p \geq \overline{S}_{p-1} + \frac{\gamma_{p-2}}{\gamma_{p-1}}\beta$. Plugging in $\overline{S}_2 \geq 0$, induction on p therefore gives

$$S^*_{\theta_p} \geq \overline{S}_p \geq \beta \cdot \sum_{i=2}^{p-1} \frac{\gamma_{i-1}}{\gamma_i} \tag{33}$$

for all $p \geq 3$. We define $\gamma_{\min} = \min\{\gamma_1, \ldots, \gamma_{p-1}\}$ and $\gamma_{\max} = \max\{\gamma_1, \ldots, \gamma_{p-1}\}$. By Lemma 4.4 and because $\eta_1 \leq \cdots \leq \eta_{p-1}$, we have $\gamma_{\min} \geq \eta_1/k$ and $\gamma_{\max} \leq \eta_p + \Delta_{\max}$. From $\alpha \geq 1$ and (24), we have $\eta_1 \geq (\alpha-1)\sigma_{\min} + \beta$. Further, we have $\eta_p = (\alpha-1)S^*_{\theta_p} + \beta < 2\beta$. We therefore have $\gamma_{\min} \geq [(\alpha-1)\sigma_{\min} + \beta]/k$ and $\gamma_{\max} < 2\beta + \Delta_{\max}$ and thus

$$\frac{\gamma_{\min}}{\gamma_{\max}} \;\geq\; \frac{(\alpha-1)\sigma_{\min} + \beta}{k(2\beta + \Delta_{\max})} \overset{(12)}{\geq} \frac{\max\{\beta, \Delta_{\max}\}}{k(2\beta + \Delta_{\max})} \;\geq\; \frac{1}{3k}.$$

The lemma now follows from (33) and from Inequality (I) of Lemma 4.7. □

On the other hand, as soon as $S^*_{\theta_p} > \max\left\{\sigma_{\min}, \frac{\beta}{\alpha-1}\right\}$, the effect of the β-term in (11) becomes relatively small. As a second case, therefore, we analyze how the service cost grows by just considering terms that depends on α (and not on β).

Lemma 4.9. *Let $p_0 \geq 2$ be a phase for which $\overline{S}_{p_0} \geq \overline{S}_{p_0-1} \geq S_0 := \max\{\sigma_{\min}, \frac{\beta}{\alpha-1}\}$. For any phase $p > p_0$, we have*

$$S^*_{\theta_p} \geq S_0 \cdot \left(1 + \frac{\sqrt{\alpha}-1}{(3k)^{\frac{1}{p-p_0}}}\right)^{p-p_0} \;\geq\; \frac{S_0}{3k} \cdot \alpha^{\frac{p-p_0}{2}}.$$

Proof. By Lemma 4.6, using $\beta \geq 0$, for all $p > p_0$, we get $\overline{S}_p \geq \left(1 + (\alpha-1)\frac{\gamma_{p-2}}{\gamma_{p-1}}\right) \cdot \overline{S}_{p-1}$. Induction on p therefore gives

$$S^*_{\theta_p} \geq \overline{S}_p \geq \overline{S}_{p_0} \cdot \prod_{i=p_0}^{p-1} \left(1 + (\alpha-1)\frac{\gamma_{i-1}}{\gamma_i}\right) \tag{34}$$

for all $p \geq p_0$. Similarly to before, we define $\gamma_{\min} = \min\{\gamma_{p_0-1}, \ldots, \gamma_{p-1}\}$ and $\gamma_{\max} = \max\{\gamma_{p_0-1}, \ldots, \gamma_{p-1}\}$. By Lemma 4.4, the assumptions regarding p_0, and because the values η_i are non-decreasing in i, we have

$$\begin{aligned} \gamma_{\min} &\geq \frac{\eta_{p_0-1}}{k} \geq \frac{\max\{(\alpha-1)\sigma_{\min} + \beta, 2\beta\}}{k} \quad \text{and} \\ \gamma_{\max} &\leq \eta_p + \Delta_{\max} \leq (\alpha-1)S^*_{\theta_p} + \beta + \Delta_{\max} \leq 3(\alpha-1)S^*_{\theta_p}. \end{aligned}$$

The last inequality follows because $S^*_{\theta_p} \geq \overline{S}_p \geq \overline{S}_{p_0} \geq \max\{\sigma_{\min}, \frac{\beta}{\alpha-1}\}$ and by applying (12). We can now apply Inequality (II) from 4.7 to obtain

$$\begin{aligned} S^*_{\theta_p} \geq \overline{S}_p &\geq \overline{S}_{p_0} \cdot \left(1 + (\alpha-1)\left(\frac{\gamma_{\min}}{\gamma_{\max}}\right)^{\frac{1}{p-p_0}}\right)^{p-p_0} \\ &\geq \overline{S}_{p_0} \cdot \left(1 + (\alpha-1)\left(\frac{\max\{(\alpha-1)\sigma_{\min} + \beta, 2\beta\}}{3k(\alpha-1)S^*_{\theta_p}}\right)^{\frac{1}{p-p_0}}\right)^{p-p_0}. \end{aligned} \tag{35}$$

In the following, assume that

$$S^*_{\theta_p} \leq \max\left\{\sigma_{\min}, \frac{\beta}{\alpha-1}\right\}\alpha^{\frac{p-p_0}{2}}. \tag{36}$$

Note that if (36) does not hold, the claim of the lemma is trivially true. By replacing $S^*_{\theta_p}$ on the right-hand side of (35) with the upper bound of (36), we obtain

$$\begin{aligned} S^*_{\theta_p} \geq \overline{S}_p &\geq \overline{S}_{p_0} \cdot \left(1 + (\alpha-1)\cdot\left(\frac{(\alpha-1)\sigma_{\min}+\beta}{3k(\alpha-1)\max\left\{\sigma_{\min}, \frac{\beta}{\alpha-1}\right\}\alpha^{\frac{p-p_0}{2}}}\right)^{\frac{1}{p-p_0}}\right)^{p-p_0} \\ &\geq \overline{S}_{p_0}\cdot\left(1+\frac{\alpha-1}{(3k)^{\frac{1}{p-p_0}}\sqrt{\alpha}}\right)^{p-p_0} \\ &\geq \overline{S}_{p_0}\cdot\left(1+\frac{\sqrt{\alpha}-1}{(3k)^{\frac{1}{p-p_0}}}\right)^{p-p_0} \geq \frac{\overline{S}_{p_0}}{3k}\cdot\alpha^{\frac{p-p_0}{2}}. \end{aligned}$$

The lemma then follows because we assumed that $\overline{S}_{p_0} \geq \max\left\{\sigma_{\min}, \frac{\beta}{\alpha-1}\right\}$. □

As a result of Observation 4.2 and Lemmas 4.8 and 4.9, it is possible to prove the main upper bound theorem as given in Section 3.

Proof of Theorem 3.1. First note that by Observation 4.2, the movement cost of our algorithm by time θ_p is at most

$$M_{\theta_p} \leq (p-1)k + 1 \leq pk. \tag{37}$$

Together with the lower bounds on $S^*_{\theta_p}$ of Lemmas 4.8 and 4.9, this allows to derive an upper bound on the movement cost of our algorithm as a function of $S^*_{\theta_p}$. Note that as all upper bound claimed in the theorem have an additive term of $O(k)$ (with no specific constant), it is sufficient to prove that the theorem holds for all time $t = \theta_p$, where $p \geq 2$ is a phase number.

Let us first consider the case where $\alpha = 1$. Because in that case $\beta/(\alpha-1)$ is unbounded, we can only apply Lemma 4.8 to upper bound the movement cost as a function of S^*_t. We choose $\ell \geq 1$ and assume that $\beta \geq k(3\rho k)^{1/\ell}$. Together with (37), for $p \geq \ell + 2$, Lemma 4.8 then gives

$$S^*_{\theta_p} \geq k(3\rho k)^{\frac{1}{\ell}}\cdot(p-2)\cdot(3k)^{-\frac{1}{\ell}} = \rho^{\frac{1}{\ell}}\cdot k(p-2) \geq \rho^{\frac{1}{\ell}} M_t - 2\rho^{\frac{1}{\ell}}k.$$

The first part of Theorem 3.1 then follows because the total movement cost for the first $\ell+2$ phases is at most $O(\ell k)$. The special cases are obtained as follows. For $\beta = \Omega(k)$, we set $\rho = \Theta(1)$ and $\ell = \Theta(\log k)$, whereas for $\beta = \Omega(k\log^{\varepsilon} k)$, we set $\rho = \Theta(1)$ and $\ell = \Theta\big(\frac{\log k}{\log\log k}\big)$.

Let us therefore move to the case where $\alpha > 1$. Let p_0 be the first phase $p_0 \geq 2$ for which $S^*_{\theta_{p_0}} \geq S_0$, where $S_0 = \max\left\{\sigma_{\min}, \frac{\beta}{\alpha-1}\right\}$ as in Lemma 4.9. Further, we set $p_1 = p_0 + \lceil 2\log_\alpha(3k)\rceil$. Using Lemma 4.9, for $p \geq p_1$, we have

$$S_{\theta_p} \geq \frac{S_0}{3k}\cdot\alpha^{\frac{p-p_1}{2}}\alpha^{\frac{p_1-p_0}{2}} \geq S_0\cdot\alpha^{\frac{p-p_1}{2}}.$$

We therefore get

$$M_{\theta_p} \leq k\cdot p \leq k\left(p_1 + 2\log_\alpha\frac{S^*_{\theta_p}}{S_0}\right) \leq k\left(p_0 + 1 + 2\log_\alpha S^*_{\theta_p} + \log_\alpha\frac{3k}{S_0}\right).$$

The second claim of Theorem 3.1 then follows by showing that

$$p_0 = O\big(\min\{\frac{\log k}{\log\log k}, \log_\alpha k\}\big).$$

If $S_0 = \sigma_{\min}$, we have $p_0 = 2$. Otherwise, we can apply Lemma 4.8 to upper bound p_0 as the smallest value p_0 for which $\frac{\beta}{\alpha-1} = \beta(p-2)(3k)^{-1/(p-2)}$. For $\alpha = O\big(\frac{\log k}{\log\log k}\big)$, the assumption that α is at least $1+\varepsilon$ for some constant $\varepsilon > 0$ gives that $p_0 = \Theta\big(\frac{\log k}{\log\log k}\big)$. Otherwise, (i.e., for large α), we obtain $p_0 = \Theta(\log_{\alpha-1} k) = \Theta(\log_\alpha k)$. □

5 Lower Bound

The aim of this section is to prove our lower bound theorem stated in Section 3. For convenience, we in the following assume that the cost for a single movement is 1. As discussed in Section 3, the lower bound even holds for a natural special case where each node $v \in V$ can only have either 0 or 1 resources (i.e., having more than 1 resources does not improve the service cost), where the service cost at a node with 1 resource is 0, the service cost at a node with 0 resources is equal to the number of demands at the node, and where an offline algorithm is also charged for its movement cost.

We start with a sketch of the lower bound proof. Assume that we are given parameters $\alpha \geq 1$ and $\beta \geq 0$ and an algorithm $\mathcal{A}$ which guarantees that at all times t, $S_t^{\mathcal{A}} < \alpha(S_t^* + M_t^*) + \beta$. Given $\mathcal{A}$, we construct an execution in which $\mathcal{A}$ has to perform a large number of movements while the optimal service cost does not grow too much. Such an execution is constructed as follows. Analogously to the analysis of the upper bound, we divide the time into phases such that in each phase, $\mathcal{A}$ has to move $\Omega(k)$ resources and the optimal service cost grows as slowly as possible. For p phases, we define a sequence of integers $k/3 \geq n_1 \geq n_2 \geq \ldots n_p \geq 1$ and values $\Gamma_1 < \Gamma_2 < \cdots < \Gamma_p$. In the following, let v be a free node if v does not have a resource. Roughly, at the beginning of a phase i, we choose a set N_i of n_i (ideally) free nodes and make sure that all these nodes have Γ_i demands. Note that constructing an execution means to determine where to add the demand in each iteration. The value Γ_i is chosen large enough such that throughout Phase i a service cost of $n_i\Gamma_i$ is sufficiently large to force an algorithm to move. Hence, whenever there are n_i free nodes with Γ_i demands, $\mathcal{A}$ has to move at least one resource to one of these nodes. For each such movement, we pick another free node that currently has less than Γ_i demands and make sure it has Γ_i demands. We proceed until there are k nodes with Γ_i demands at which point the main part of the phase ends. Except for the nodes in N_i, each of the k nodes with Γ_i demands leads to a movement of $\mathcal{A}$ and therefore, $\mathcal{A}$ has to move at least $k - n_i = \Omega(k)$ resources in Phase i. At the end of Phase i, we can guarantee that there are exactly k nodes with Γ_i demands, n_i nodes with Γ_{i-1} demands, $n_{i-1} - n_i$ nodes with Γ_{i-2} demands, etc. Assuming that for all v, $\sigma_v(x,y) = (1-x)y$, we can then compute the optimal service cost after Phase p as $n_p\Gamma_{p-1} + \sum_{i=3}^{p}(n_{i-1} - n_i)\Gamma_{i-2}$. By choosing the values n_p appropriately, we obtain the claimed bounds. In the following, we provice a formal proof of Theorem 3.3.

Proof of Theorem 3.3. We need n to be sufficiently large, for simplicity assume that $n \geq 3k$. As we can assume that each node either has 0 or 1 resource, we slightly overload notation and simply denote a feasible solution by a set $F \subset V$ of size $|F| = k$. For convenience, we further define $V' := V \setminus F_0$ to be the nodes with no resources at the beginning. As a first step in our execution, we pick a set $N_0 \subset V'$ of $|N_0| = k$ nodes and we add 2 demands to all nodes in N_0. Note that an optimal offline algorithm would now move a resource to each node in N_0 resulting in a solution with service cost 0 and movement cost k. For the rest of the execution, we only add demands to nodes in V'. Hence, from now on, at each point t in the execution, an optimal solution F_t^* places resources at the k nodes in V' with the most demands (breaking ties arbitrarily if there are several nodes with the same number of demands). Also, at a time t after adding the 2 demands

to each node in N_0, the optimal movement cost will remain to be k and the optimal service cost is equal to the total number of demands at nodes in $V' \setminus F_t^*$ for an arbitrary optimal solution F_t^*.

From now on time is divided into phases. As described in the outline, we define integers $\Gamma_1 < \Gamma_2 < \dots$ such that at the end of Phase i, there are exactly k nodes with Γ_i demands (and all other nodes have fewer demands). For each phase i, we define V_i to be this set of nodes with k demands and for convenience, we also define $V_0 := N_0$ and $\Gamma_0 := 2$, i.e., we interpret the above execution prefix as Phase 0. We also fix integers $k/3 \geq n_1 \geq n_2 \geq \dots \geq 1$ and at the beginning of each phase i, we pick a set N_i of n_i nodes to which we directly add demands so that all of them have exactly Γ_i demands. For $i = 1$, we pick N_1 as an arbitrary subset of $V' \setminus V_0$. For $i \geq 2$, we choose N_i as an arbitrary subset of $V_{i-2} \setminus V_{i-1}$. Clearly, at the end of Phase i, we have $N_i \subseteq V_i$ as otherwise there would be more than k nodes with exactly Γ_i demands. Note that because $N_{i-1} \subseteq V_{i-1}$ and because $N_{i-1} \cap V_{i-2} = \emptyset$, $V_{i-2} \setminus V_{i-1}$ contains $n_{i-1} \geq n_i$ nodes and it is therefore possible to choose N_i as described. Note also that because $N_i \subset V_{i-2} \setminus V_{i-1}$, at the beginning of Phase i all nodes in N_i have exactly Γ_{i-2} demands. The remaining ones of the k nodes that end up in V_i (and thus have Γ_i demands at the end of Phase i) are chosen among the nodes in V_{i-1}. Consequently, at the end of Phase $i-1$ and thus at the beginning of Phase i, there are exactly k nodes V_{i-1} with Γ_{i-1} demands, n_{i-1} nodes $V_{i-2} \setminus V_{i-1}$ with Γ_{i-2} demands, $n_{i-2} - n_{i-1}$ demands $V_{i-3} \setminus (V_{i-2} \cup N_{i-1})$ with Γ_{i-3} demands, $n_{i-3} - n_{i-2}$ nodes with Γ_{i-4} demands, and so on. Now, n_i of the nodes in $V_{i-2} \setminus V_{i-1}$ are chosen as set N_i and we increase their number of demands to Γ_i. From now on, throughout Phase i, there are $k + n_i$ nodes with at least Γ_{i-1} demands such that at most k of these nodes have Γ_i demands. The number of nodes with less than Γ_{i-1} demands is the same as at the end of Phase $i-1$. In fact nodes that are not in $V_{i-1} \cup N_i$ do not change their number of demands after Phase $i-1$. As a consequence, after increasing the number of demands in N_i to Γ_i, the optimal service cost remains constant throughout Phase $i \geq 1$ and it can be evaluated to

$$\Sigma_i^* := n_i \cdot \Gamma_{i-1} + \sum_{j=1}^{i-1} (n_j - n_{j+1}) \Gamma_{j-1}.$$

For convenience, we also define $\Sigma_0^* := 0$.

Now assume that there is an algorithm $\mathcal{A}$ which guarantees that $S_t^{\mathcal{A}} < \alpha(S_t^* + M_t^*) + \beta$ at all times $t \geq 0$. When Phase 1 starts, we have $M_t^* = k$ and therefore the condition simplifies to $S_t^{\mathcal{A}} < \alpha S_t^* + \beta'$, where $\beta' = \beta + \alpha k$.

In the following, let v be a free node at some point in the execution, if the algorithm currently has no resource at node v. We now fix a phase $p \geq 1$ and assume that we are at a time t, when we have already picked the set N_p and increased the number of demands of nodes in N_p to Γ_p. By the above observation, we have $S_t^* = \Sigma_p^*$ and therefore $\mathcal{A}$ is forced to move if there are n_p free nodes with Γ_p demands and if we choose Γ_p such that

$$\gamma_p := \Gamma_p - \Gamma_{p-1} \geq \frac{(\alpha - 1)\Sigma_p^* + \beta'}{n_p}. \tag{38}$$

We can now describe how and when the remaining $k - n_p$ nodes of V_p are chosen after picking the nodes in N_p. As described above, the nodes are chosen from $V_{p-2} \setminus V_{p-1}$. We choose the nodes sequentially. Whenever we choose a new node from V_p, we pick some free node $v \in V_{p-1}$ with less than Γ_p demands and increase the number of demands of v to Γ_p. As described above, Γ_p is chosen large enough (as given in (38)) such that throughout Phase p there are never more than $n_p - 1$ free nodes with Γ_p demands. Because $|N_p \cup V_{p-1}| = k + n_p$, as long as there are at most k nodes with Γ_p demands there always needs to be a free node $v \in V_{p-1}$ that we can pick and we actually manage to add k nodes to V_p. Note that the choice of which nodes to add to V_p depends on the actions of the algorithm and consequently in the case of a randomized algorithm $\mathcal{A}$, our lower bound only holds for an adaptive online adversary.

Let us next count the number of movements of $\mathcal{A}$ in a given phase p. At each point in time t during the phase, let Φ_t be the number of free nodes with Γ_p demands (possibly including a node v that we already chose to be added to V_p). We know that for all t, $\Phi_t < n_p$. Whenever we decide to add a new node v to V_p, Φ_t increases by 1 (as v is a free node). The value of Φ_t can only decrease when $\mathcal{A}$ moves a resource and each resource movement reduces the value of Φ_t by at most 1. As after fixing N_p, we add $k - n_p$ nodes to V_p, we need at least $k - 2n_p \geq k/3$ movements to keep Φ_t below n_p throughout the phase. Consequently, every online algorithm $\mathcal{A}$ has to do at least $k/3$ movements per phase.

To finish the proof, it remains to upper bound the optimal offline cost Σ_p^* as a function of α, β, and p. Using (38), for all $p \geq 0$, we have

$$\Sigma_p^* = \sum_{i=1}^{p} n_p \cdot \gamma_{p-1}$$

For $p \geq 1$, we then get

$$\begin{aligned} \Sigma_p^* &= \frac{n_p}{n_{p-1}}\big((\alpha-1)\Sigma_{p-1}^* + \beta'\big) + \Sigma_{p-1}^* \\ &= \left(1 + (\alpha-1)\frac{n_p}{n_{p-1}}\right) \cdot \Sigma_{p-1}^* + \beta' \cdot \frac{n_p}{n_{p-1}}. \end{aligned} \tag{39}$$

In the following, we for simplicity assume that for $i = 1, 2, \ldots, p$, n_i does not have to be an integers. For integer n_i, the proof works in the same way, but becomes more technical and harder to read. We fix the values of n_i as

$$n_i := (k/3)^{\frac{p-i}{p-1}}$$

such that $n_1 = k/3$ and $n_p = 1$. For all $i \geq 1$, we then have $\frac{n_i}{n_i - 1} = (\frac{k}{3})^{-\frac{1}{p-1}}$. Equation (39) now simplifies to

$$\Sigma_p^* = \left(1 + \frac{\alpha - 1}{(k/3)^{1/(p-1)}}\right) \cdot \Sigma_{p-1}^* + \beta' \cdot \frac{1}{(k/3)^{1/(p-1)}}.$$

The claim of the theorem now follows analogously to the argument in the proof of Theorem 3.1. □

6 Conclusions

In the present paper, we considered the following online problem: There is a set of k moveable resources, each located at one of n nodes. Demands arrive in an online fashion at the nodes and the resources have to be moved between nodes in order to keep the service cost (modeled using a general cost function for each node) low. We give nearly tight upper and lower bounds for the problem of minimizing the total number of movements and the overall service cost as a function of the current optimal offline cost. As a special case, our bounds apply to an online version of the standard mobile metric facility location problem [1, 14] if the underlying metric is uniform (i.e., all distances are equal). We hope that some of the techniques of our paper are also useful when studying movement minimization for general mobile metric facility location or for similar related online problems.

References

[1] S. Ahmadian, Z. Friggstad, and C. Swamy. Local-search based approximation algorithms for mobile facility location problems. In *Proc. 24th Symp. on Discrete Algorithms (SODA)*, pages 1607–1621, 2013.

[2] A. Anagnostopoulos, R. Bent, E. Upfal, and P. V. Hentenryck. A simple and deterministic competitive algorithm for online facility location. *J. Information and Computation*, 194(2):175–202, 2004.

[3] V. Arya, N. Garg, R. Khandekar, A. Meyerson, K. Munagala, and V. Pandit. Local search heuristics for k-median and facility location problems. *SIAM J. on Computing*, 33(3):544–562, 2004.

[4] N. Bansal, N. Buchbinder, A. Madry, et al. A polylogarithmic-competitive algorithm for the k-server problem. In *Proc. 52nd Symp. on Foundations of Comp. Sc. (FOCS)*, pages 267–276, 2011.

[5] J. Byrka and K. Aardal. An optimal bifactor approximation algorithm for the metric uncapacitated facility location problem. *SIAM J. on Computing*, 39(6):2212–2231, 2010.

[6] M. Charikar, C. Chekuri, T. Feder, and R. Motwani. Incremental clustering and dynamic information retrieval. In *Proc. 29th Symp. on Theory of Computing (STOC)*, pages 626–635, 1997.

[7] E. D. Demaine, M. Hajiaghayi, H. Mahini, A. S. Sayedi-Roshkhar, S. Oveisgharan, and M. Zadimoghaddam. Minimizing movement. *ACM Trans. on Alg. (TALG)*, 5(3):30, 2009.

[8] G. Divéki and C. Imreh. Online facility location with facility movements. *Central European J. of Operations Research*, 19(2):191–200, 2011.

[9] Z. Drezner and H. W. Hamacher. *Facility location: applications and theory*. Springer, 2004.

[10] D. Fotakis. Incremental algorithms for facility location and k-median. In *Proc. 12th Annual European Symp. (ESA)*, pages 347–358, 2004.

[11] D. Fotakis. Memoryless facility location in one pass. In *Proc. 23rd Annual Symp. on Theoretical Asp. of Computer Science (STACS)*, pages 608–620, 2006.

[12] D. Fotakis. On the competitive ratio for online facility location. *Algorithmica*, 50(1):1–57, 2008.

[13] D. Fotakis. Online and incremental algorithms for facility location. *ACM SIGACT News*, 42(1):97–131, 2011.

[14] Z. Friggstad and M. R. Salavatipour. Minimizing movement in mobile facility location problems. *ACM Trans. on Alg. (TALG)*, 7(3):28, 2011.

[15] S. Guha and S. Khuller. Greedy strikes back: Improved facility location algorithms. In *Proc. 9th Symp. on Discrete Algorithms (SODA)*, pages 649–657, 1998.

[16] M. T. Hajiaghayi, M. Mahdian, and V. S. Mirrokni. The facility location problem with general cost functions. *Networks*, 42(1):42–47, 2003.

[17] K. Jain, M. Mahdian, E. Markakis, A. Saberi, and V. V. Vazirani. Greedy facility location algorithms analyzed using dual fitting with factor-revealing lp. *J. of the ACM (JACM)*, 50(6):795–824, 2003.

[18] E. Koutsoupias and C. H. Papadimitriou. On the k-server conjecture. *Journal of the ACM (JACM)*, 42(5):971–983, 1995.

[19] M. S. Manasse, L. A. McGeoch, and D. D. Sleator. Competitive algorithms for server problems. *Journal of Algorithms*, 11(2):208–230, 1990.

[20] A. Meyerson. Online facility location. In *Proc. 42nd Symp. on Foundations of Computer Science (FOCS)*, pages 426–431, 2001.

A Movement Cost of the Offline Algorithm

Even though our analysis does not explicitly deal with the movement cost of an offline algorithm, the possibility to deal with general cost functions gives an elegant way to incorporate offline movement cost into the analysis. Assume that we are given a cost function $\sigma_v(x, y)$ that does not account for any movements an offline algorithm has to do. Because an offline algorithm can defer all movements until the very end, the number of movements of such an algorithm only depends on the difference between the initial resource configuration $F_0 = \{(v, f_{v,0}) : v \in V\}$ and the solution $F = \{(v, f_v) : v \in V\}$ at the end of the algorithm. In fact, the movement cost is exactly $\chi(F_0, F)$. We can define a new cost function $\sigma'_v(x, y)$ as follows:

$$\forall v \in V, \forall x \in \{0, \ldots, k\}, \forall y \in \mathbb{N}_0 : \sigma'_v(x, y) := \sigma_v(x, y) + \frac{|x - f_{v,0}|}{2}.$$

Using (4), for any feasible solution F, we then have $S'_t(F) = S_t(F) + \chi(F_0, F)$, where $S'_t(F)$ refers to the total service cost w.r.t. the new cost function σ'_v. Hence, $S_t(F)$ exactly measures the sum of service cost and movement cost of a solution F. Of course now, in all our results, S^*_t corresponds to the combination of service and movement cost of an optimal offline algorithm. Also note that the cost function σ'_v directly implies that $\sigma_{\min} \geq 1$.